

CLAIMS

WHAT IS CLAIMED IS:

1. A hierarchical traffic management system comprising at least one traffic management node comprising:

5 a classifier operable to identify and classify incoming traffic streams; and

 a queuing system comprising a plurality of queues and operable to apply scheduling policies to said traffic streams, said plurality of queues of the queuing system each comprising:

 enqueue attributes configured to control a depth of the queue; and

10 dequeue attributes configured to control scheduling of the queue, said dequeue attributes comprising minimum bandwidth, maximum bandwidth, excess bandwidth, and priority, wherein each of said plurality of queues has one or more of said dequeue attributes defined.

15 2. The system of claim 1 wherein the packet scheduling system comprises at least three layers of hierarchy.

3. The system of claim 1 wherein at least a portion of said plurality of queues have said minimum bandwidth attribute defined.

5 4. The system claim 3 wherein at least a portion of said plurality of queues have said excess bandwidth attribute defined.

5. The system of claim 1 wherein at least a portion of said plurality of queues have said maximum bandwidth attribute defined.

10 6. The system of claim 1 wherein the attributes not defined at each of said plurality of queues have a default value.

7. The system of claim 1 wherein each of said dequeue attributes are defined.

15 8. The system of claim 1 wherein each of said plurality of queues set as a priority queue sends traffic ahead of other queues sharing a parent node.

9. The system of claim 1 wherein each of said plurality of queues having a defined priority attribute have defined a level of priority.

10. The system of claim 1 wherein priority nodes are configured as conditional or unconditional priority nodes.

11. The system of claim 1 wherein each of said plurality of queues is configured to allow for an oversubscription mode in which minimum rates are oversubscribed.

12. The system of claim 1 wherein the traffic management node further comprises a pre-queuing operator configured to operate on at least some of said incoming traffic streams before the streams enter the queuing system.

13. The system of claim 1 wherein the traffic management node further comprises a post-queuing operator configured to operate on at least some of said incoming traffic streams after the streams pass through the queuing system.

14. The system of claim 13 wherein the post-queuing operator is configured to compress packets.

15. The system of claim 1 wherein each of said plurality of queues comprises a minimum-rate propagation attribute.

16. The system of claim 1 wherein each of said plurality of queues comprises a priority propagation attribute.

17. The system of claim 1 wherein the traffic management nodes operate according to a specified behavioral model and a plurality of components within a network system comprise one or more of the traffic management nodes.

18. The system of claim 17 wherein a common configuration language is provided for user interface with the behavioral model.

19. The system of claim 1 wherein the depth of a queue is controlled by a specified maximum queue depth.

20. The system of claim 1 wherein the depth of a queue is controlled by specification of a Random Early Detection profile.

21. A method for hierarchical traffic management at a network device having a queuing system comprising a plurality of layers of hierarchy, each layer of the hierarchy configured for supporting one or more priority nodes, one or more nodes having a guaranteed minimum rate, one or more nodes designated for receiving excess
5 bandwidth, and one or more nodes having a defined maximum rate, the method comprising:

classifying incoming traffic streams; and

applying scheduling policies to said traffic streams at one or more queues, said scheduling policies comprising minimum bandwidth, maximum bandwidth, excess
10 bandwidth, and priority, wherein traffic up to a specified bandwidth is defined as priority traffic.

22. The method of claim 21 wherein said one or more of the scheduling policies at said one or more of the queues have default values applied.

23. The method of claim 21 wherein said scheduling policies further comprise priority propagation which specifies whether or not priority service at a queue is propagated through a hierarchy of the queue.

24. The method of claim 23 wherein a burst tolerance parameter is associated with a stream enabled with priority propagation.

25. The method of claim 24 wherein the burst tolerance parameter is provided
5 for each layer of hierarchy through which priority behavior propagates.

26. The method of claim 24 wherein the burst tolerance parameter indicates how much the stream may burst beyond a rate constraint of an ancestral node before the stream becomes constrained by the rate.
10

27. The method of claim 21 wherein said scheduling policies further comprise minimum rate propagation which specifies whether or not a minimum rate at a queue is propagated through a hierarchy of the queue.

15 28. The method of claim 21 further comprising enabling an oversubscription mode in which oversubscribed streams are reduced in proportion to a specified oversubscription minimum rate.

29. A computer program product method for hierarchical traffic management at a network device having a queuing system comprising a plurality of layers of hierarchy, each layer of the hierarchy configured for supporting one or more priority nodes, one or more nodes having a guaranteed minimum rate, one or more nodes designated for receiving excess bandwidth, and one or more nodes having a defined maximum rate, the method comprising:

code that classifies incoming traffic streams;

code that applies scheduling policies to said traffic streams at one or more queues, said scheduling policies comprising minimum bandwidth, maximum bandwidth, excess bandwidth, and priority, wherein traffic up to a specified bandwidth is defined as priority traffic; and

a computer-readable storage medium for storing said codes.

30. A system for hierarchical traffic management at a network device having a queuing system comprising a plurality of layers of hierarchy, each layer of the hierarchy configured for supporting one or more priority nodes, one or more nodes having a guaranteed minimum rate, one or more nodes designated for receiving excess bandwidth, and one or more nodes having a defined maximum rate, the method comprising:

means for classifying incoming traffic streams; and

means for applying scheduling policies to said traffic streams at one or more queues, said scheduling policies comprising minimum bandwidth, maximum bandwidth, excess bandwidth, and priority, wherein traffic up to a specified bandwidth is defined as priority traffic.

5

10